



THE HOME COMPUTER COURSE

4

MASTERING YOUR HOME COMPUTER IN 24 WEEKS



061 Word Processing
064 Questions and Answers
065 In-Car Computers
066 Software Hierarchies
068 Logical Gates
070 Lynx
072 Expert Systems
074 Printers
077 Basic Programming
079 Binary Arithmetic

An ©RBIS Publication

IR £1 Aus \$1.95 NZ \$2.25 SA R1.95 Sing \$4.50 USA & Can \$1.95

CONTENTS

Hardware



Micros On The Move Microcomputers will be found under the bonnet and behind the dashboard of tomorrow's family car 65

Lynx The computer of the week with a sympathetic keyboard, elegant styling and bright graphics 70

Software



Word Perfect A micro takes the frustration out of typing and can produce perfect copies every time 61

Consulting The Chip The next time you visit a specialist you could be faced with a micro 72

When Minus Equals Plus Part Two of our course in Binary 79

Basic Programming



Routine Matters We dive deeper into the clear waters of programming in Basic 77

Feedback



Questions And Answers Apples still grow in Silicon Valley and the computer has caught up with the police 64

Insights



Message Understood Through English to Basic to the electronic language of the microprocessor 66

Fit To Print Caxton never dreamed of the ink jet printer, the daisy wheel and the dot matrix printer 74

Passwords To Computing



True Or False? We open the gateways to logic and discover how computers make 'rational' decisions 68

Next Week

• We review the BBC Micro, a best-selling computer that will soon be able to receive programs broadcast directly to your television screen

• We take a look at the fascinating history of computers, from the first 19th-century prototype to the giant machines that are vital to the space race

• Computers in the classroom - we discover how educational programs are developed into effective aids to all aspects of learning



YOUR BINDER
ORDER FORM
IS WITH
THIS
ISSUE.



Editor Gareth Jefferson; **Art Editor** David Whelan; **Production Editor** Catherine Cardwell; **Picture Editor** Claudia Zeff; **Designer** Hazel Bennington; **Art Assistants** Steve Broadhurst, Julie-Anne Chambers, Liz Dixon; **Sub Editor** Teresa Bushell; **Researcher** Melanie Davis; **Consultant Editor** David Tebbutt; **Project Manager** Ram Egerton; **Contributors** Tim Heath, Peter Jackson, Ray Hammond, Garry Marshall, Rom Smith, Ian White; **Group Art Director** Perry Neville; **Managing Director** Stephen England; **Editorial Director** Brian Innes; **Project Development** Peter Brooksmith; **Executive Editor** Chris Cooper; **Production Co-ordinator** Ian Paton; **Circulation Director** David Breed; **Marketing Director** Michael Joyce; **Designed and produced by** Bunch Partworks Ltd; **Published by** Orbis Publishing Ltd; © 1983 by Orbis Publishing Ltd; **Typeset by** Universe; **Reproduction by** Mullis Morgan Ltd; **Printed in Great Britain by** Artisan Press Ltd, Leicester

HOME COMPUTER COURSE - Price UK 80p IR £1.00 AUS \$1.95 NZ \$2.25 SA R1.95 SINGAPORE \$4.50 USA and CANADA \$1.95

How to obtain your copies of HOME COMPUTER COURSE - Copies are obtainable by placing a regular order at your newsagent.

Back Numbers UK and Eire - Back numbers are obtainable from your newsagent or from HOME COMPUTER COURSE. Back numbers, Orbis Publishing Limited, 20/22 Bedfordbury, LONDON WC2N 4BT at cover price. AUSTRALIA: Back numbers are obtainable from HOME COMPUTER COURSE. Back numbers, Gordon & Gotch (Aus) Ltd, 114 William Street, PO Box 7676 Melbourne, Vic 3001. SOUTH AFRICA, NEW ZEALAND, EUROPE & MALTA: Back numbers are available at cover price from your newsagent. In case of difficulty write to the address in your country given for binders. South African readers should add sales tax.

How to obtain binders for HOME COMPUTER COURSE - UK and Eire: Details of how to obtain your binders (and of our special offer) will be in issue 7. EUROPE: Write with remittance of £5.00 per binder (incl. p&p) payable to Orbis Publishing Limited, 20/22 Bedfordbury, LONDON WC2N 4BT. MALTA: Binders are obtainable through your local newsagent price £3.95. In case of difficulty write to HOME COMPUTER COURSE BINDERS, Millier (Malta) Ltd, M. A. Vassalli Street, Valletta, Malta. AUSTRALIA: For details of how to obtain your binders see inserts in early issues or write to HOME COMPUTER COURSE BINDERS, First Post Pty Ltd, 23 Chandos Street, St. Leonards, NSW 2065. The binders supplied are those illustrated in the magazine. NEW ZEALAND: Binders are available through your local newsagent or from HOME COMPUTER COURSE BINDERS, Gordon & Gotch (NZ) Ltd, PO Box 1595, Wellington. SOUTH AFRICA: Binders are available through any branch of Central Newsagency. In case of difficulty write to HOME COMPUTER COURSE BINDERS, Intermap, PO Box 57394, Springfield 2137.

Note - Binders and back numbers are obtainable subject to availability of stocks. Whilst every attempt is made to keep the price of the issues and binders constant, the publishers reserve the right to increase the stated prices at any time when circumstances dictate. Binders depicted in this publication are those produced for the UK market only and may not necessarily be identical to binders produced for sale outside the UK. Binders and issues may be subject to import duty and/or local taxes, which are not included in the above prices unless stated.



IAN MCKINWELL

Portable Processor

Word processing is rapidly becoming one of the most popular applications for microcomputers. New machines are being designed with features specifically geared to this function. These include 80 column screens (to display the full width of a typed letter), built-in disk drives (sometimes a word processing disk is even included in the price) and programmable function keys which are used to manipulate the text. Some machines, like the Ajile pictured here, are designed to be portable — ideal for journalists or jet-setting executives!

Word Perfect

With the right software, your micro is transformed into a word processor, editing, organising and storing your words and even correcting your spelling mistakes

Word processing is one of the most useful tasks that can be performed on a microcomputer. But the phrase 'word processing' has hindered public understanding of this powerful tool. 'How can you process words?' is the usual reaction.

Until very recently tube and bus advertisements offered 'word processors for your office'. What these ads neglected to say was that their expensive hardware was no more than a microcomputer specially adapted to run word processing programs. Word processors are less flexible than the ordinary microcomputer, because they can only perform one task.

Perhaps the kind of program that the phrase 'word processing' describes should have been called 'computer-assisted typing'. With the

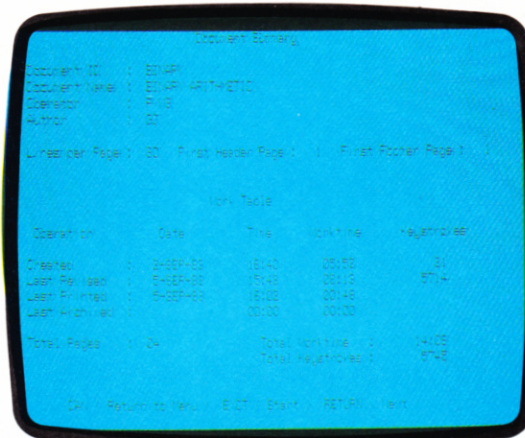
addition of a printer, most home computers can now run some form of word processing or text editing program. But it is not until computer owners try word processing for themselves that they realise how useful such a program is.

When used as a word processor, a computer displays words on the screen as they are typed in, just as a typewriter prints them on paper. The larger microcomputers can display 80 characters across the screen, representing your 'page'. On smaller computers the user requires more patience. He has to put up with a much narrower screen and on some models has to allow for the lack of lower case (small) letters. The user also has to remember that the smaller machine can only store a limited amount of text.

The program offers several sophisticated aids. All word processing programs sense the end of each line as it approaches and automatically 'wrap round', carrying the whole of the last word down to the beginning of the next line. This means that the typist no longer has to be concerned about 'carriage return' at the end of each line. Instead, he or she can type in an endless stream, while the program creates each new line as necessary. Where a new paragraph is to be started, however, the typist has to press the RETURN key.

On a conventional typewriter, you are faced with the prospect of correcting a mistake mechanically, usually by whiting out the error and typing over the top. It is pretty messy. If there are one or two corrections, the only alternatives are to send out a messy letter or to start again. With word processing, the problem is solved. The flashing cursor on the screen is, as always, the indicator of your current position. You move it back along the already written words until it is at the point where, say, you have typed the wrong word. You are then able to make the mistake disappear and type in the correction.

Once the editing power of a word processing program is realised, users are motivated to be more thoughtful in the preparation of the text. For example, it is possible to use the INSERT command to add a word, a whole sentence, or a paragraph just as easily as a single letter. This encourages the user to reconsider what is being said in the letter or

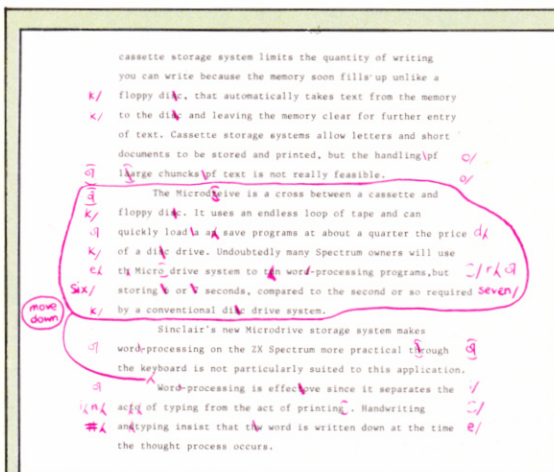


Seeing The Menu

The photograph shows the 'menu' from a sophisticated word processing package. The menu will appear on the screen as soon as you insert the software and guides you to the various editing functions available on the word processor. Examples of editing functions are: tab setting and margin positions, line spacing, counting the number of words in the document, rearranging paragraphs and creating an index

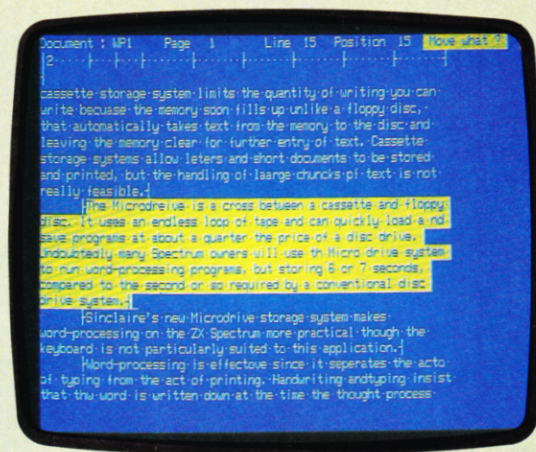
some degree of word processing power. The Sinclair ZX81 can operate with a simple text-editing program, which allows the user to write a letter or document on the screen and then edit out corrections. Text editing is the phrase usually applied to a limited word processing program, which may be able to handle a page or two of text but which is unable to manipulate and store longer documents. The small RAM in the ZX81 and similarly-sized computers severely limits the amount of text that can be displayed and worked on.

One problem with the ZX81 is the touch-sensitive keyboard, which prohibits any real



Electronic Editing

The main advantage of a computer-based word processor system over a typewriter is its amazing flexibility and speed. Whereas the typewritten text entails laborious editing and correcting, with a word processor it can all be done at electronic speed. And all these operations can be done while the text is displayed on the screen. A sophisticated system can: search through the text;



change words; transpose lines from one position to another; make corrections automatically throughout the text; and even correct your grammar and spelling. In fact, the word processor can save so much time that it is now gradually beginning to replace the ordinary typewriter.

Compare the 'messy' corrections on the typescript to the clean electronic editing of the word processor

document. The instruction to delete text is also easy. A command causes unwanted words and letters simply to disappear off the screen and remaining text to close up, restoring the perfect appearance of the page. Many professional authors and journalists are now using word processors and they generally report that both the quality and quantity of their work has improved.

Even the smallest microcomputers can offer

attempt at fast typing. Although the keyboard on the ZX Spectrum is considerably better, it is still not the mechanical kind that typists are used to. If you are considering buying a micro for word processing use, it is a good idea to examine the keyboards of the machines under consideration as these have a significant effect on typing speed and comfort (see page 32).

A computer of 16 Kbytes or 32 Kbytes,

however, is capable of quite useful word processing. After the purchase of a printer, the main hardware problem becomes one of having enough memory in which to store your text. It is possible to run a word processing program using a cassette for storage. However, a cassette storage system limits the amount of text you can write because the memory quickly fills up. This is unlike a floppy disk that automatically takes text from the memory leaving it clear for further entry of text. Cassette storage allows letters and short documents to be stored and printed, but the handling of large chunks of text isn't really feasible.

Sinclair's new Microdrive storage system makes word processing on the ZX Spectrum more practical. The Microdrive is a cross between a

of a letter or a poem. (See our feature on printers on page 74.)

Manufacturers are striving to produce inexpensive word processing systems and some chip programs are now becoming available for home computers. These store word processing programs' instructions on a chip that can be plugged into the computer's circuit board. They are very useful when there is no disk drive unit available. The advantage is that the program may be loaded quickly and is ready for instant use without waiting for it to load via a cassette tape or disk. If the computer's RAM is large enough (from 32 Kbytes upwards), you will be able to create a document of up to perhaps 5,000 words in length and edit it to your heart's content.

MICROCOMPUTER	PRICE	SOFTWARE	SUPPLIED ON CASSETTE		SUPPLIED ON DISK		TYPEWRITER STYLE KEYBOARD	PRICE
Sinclair Spectrum	£125	Wordprocessor	●					£9.95
Commodore Vic 20	£130	Vicwriter	●			●		£19.95
Dragon 32	£175	Gemini	●			●		£19.95
New Brain	£269	Propen	●					£40.25
Commodore 64	£299	Visawrite			●	●		£79
Atari 800	£300	Keyword			●	●		£115
BBC Micro	£399	View			●	●		£59.80
Epson HX20	£486	Deckmaster	●			●		£34.50

The Price Of Word Processing

The table above shows a selection of the word processing packages available for a range of microcomputers at both ends of the market. If you plan to use your word processor for long periods of time, make sure that you can type comfortably on your computer's keyboard. Calculator or membrane-style keyboards were developed to reduce the cost of manufacture and are best suited for games playing and writing short programs. It is also an advantage if your computer has the special programmable function keys. These are often employed

on the more sophisticated word processing packages and reduce the number of commands you need to type into your keyboard. The prices quoted are typical of what you can expect to pay. Obviously the purchase of a printer is essential if you want a printed copy. Printers vary tremendously in the way of print quality and speed of execution (see page 74). Ensure that your word processor and printer are compatible for the various tasks you wish them to perform. If you wish to write letters, a reasonably-priced word processor and printer would be adequate. But for lengthy text, a more expensive combination would be necessary.

cassette and a floppy disk. It uses an endless loop of tape and can quickly load and save programs at about a quarter of the price of a disk drive. Undoubtedly many Spectrum owners will use the Microdrive system to run word processing programs, but storing and recalling pages is likely to take six or seven seconds, compared with the one second or so required by a conventional disk drive system.

Word processing is effective because it separates the act of composition from the act of printing. Both handwriting and typing demand that the word is written down at the same time as the thought process occurs. With word processing, no words appear on paper until the composition on the screen is right. But getting the words onto paper demands a printer, and the cheap printer that is adequate for listing computer programs is unlikely to be able to produce a satisfactory copy

If you wish to store the edited text after you have printed it, you will need to save it onto a cassette tape; a process that will take some minutes. The word processing chip is unable to store the text you create. If you want to write your novel with a word processing program, you need to know how well the program can handle very large chunks of text.

Some sophisticated word processing programs can perform useful extra functions. The automatic dictionary, or spelling checker, is among the most popular inclusions. For this facility you really need a disk system. The dictionary checks the words contained in a document against its own stored words. It indicates the words it doesn't recognise, and invites the user to correct them.

As the use of word processing spreads, it is likely to be regarded as an essential office skill and an ideal device for correspondence.

Questions And Answers

What is the 'fifth generation' of computers, and what were the first four?

The fifth generation is the stage of computer development towards which computer engineers and programmers are currently working. Such computers are expected to represent a significant leap beyond present technology.

The term 'fifth generation' was adopted by the Japanese to describe their long term research projects. Computers of the fifth generation will not have keyboards, and they will not have to be programmed in computer languages such as BASIC and PASCAL. Instead we can expect to be able to talk to our computers and they will answer back in whatever language we require. These fifth generation computers are also likely to be able to write their own programs as solutions to problems we present them with.

The first generation of computers were the first all-electronic computing devices, invented just after the Second World War. The second generation used essentially the same concepts with transistors instead of valves and relays, making computers smaller and cheaper. The third generation used integrated circuits (electronic circuits mounted on silicon). These integrated circuits were the earliest form of the microchip and marked the beginning of the reduction in price of computers. However these computers were still too expensive for the average home or office.

The fourth generation represents current technology. These

computers use LSI (Large Scale Integration) circuits. The development of these microchips has brought the computer within the range of most people's budgets.



Where is Silicon Valley?

Silicon Valley is the name given to an area of land around San Jose to the south of San Francisco, where the headquarters or research departments of most of the large American computer and microelectronics companies are situated. The reason why so much expertise is collected together in such a small area is purely historical — there are no natural resources beneficial to the manufacture of microchips! Until 20 years ago the area was known only for producing fruit.



Is all the loose 'untidy' wiring at the back of some computers necessary?

Most of today's microcomputers are designed with as much

thought put into the outward appearance as that of the electronics. 'Untidy' wires are normally hidden away. But with some advanced research computers, the loose wiring is very important. Electricity moves at the speed of light, but it still takes a certain time to travel down a wire. These research computers work so quickly that the information has to arrive at the right place at exactly the right time. The lengths of the wires are calculated precisely to ensure the timing is perfect.



Computers are often advertised as having a Z80 or 6502 microprocessor.

What is the significance of these numbers?

The numbers themselves have no significance — '6502' is merely the identifying reference or 'name' for a particular microprocessor chip; 'Z80' is another. All computers which are based on the same microprocessor understand the same set of fundamental instructions (called Machine Code) from which programs are built up. However, programs are usually written by the user in a high level language such as BASIC and then interpreted into machine code by the computer. So unless you specifically want to write programs directly in machine

code, it doesn't make any difference what sort of microprocessor your computer has.

Though some types of microprocessor operate at a higher speed than others, the rate at which you see things happen in a typical application is far more dependent on the way in which the software has been written.



How can computers help in the fight against crime?

The Police National Computer Unit was set up in 1968 and the first file of information (stolen and suspect vehicles) was installed in 1974. Since this time, data has been included to list the details of all people with criminal records. The police are currently exploring the possibility of using microcomputers in 'incident rooms' to cope with emergencies. At present a few police cars have computer terminals through which they can access information from the central computer. The current state of storage technology means that it would be feasible for the police to store the personal details of every citizen. Fortunately, the government has implemented various 'watchdog' committees to safeguard the interests of the innocent.



Micro On The Move

No more peering at maps or worrying about running out of petrol: the car of the future will see you safely and economically to your destination



Dynamic Dashboard

The dashboard of your next car could well look similar to the one pictured here. All the moving dials and needles used in conventional dashboards have been replaced by a computer-controlled read-out terminal. The light-emitting

diodes can monitor speed, the petrol level, and the temperature with greater speed and accuracy than electro-mechanical instruments. Dashboards such as these are already in use on current models of cars such as the Austin Maestro

One of the ways in which we will see the direct relevance of computer technology at work is in the car. Already you can buy a vehicle that will calculate when it needs a service and tell you when a visit to the garage is due.

The car is able to do this because sensors attached to various engine points feed details such as mileage and temperatures into a micro-processor, which evaluates the conditions under which the car has been driven. The driver knows when a service is due because the tiny computer operates a series of green, amber and red lights on the dashboard. When the green lights go out and the red lights come on, the driver knows the car needs a service.

There are cars made in Europe and Japan that speak to the driver, telling him to put on the seatbelt or warning him of impending engine trouble such as overheating or low water and oil levels. The car can do this because it has a voice synthesiser; a computer into which the characteristics of a human voice have been programmed digitally. This is called digital sound because the sound waves are turned into the binary numbers that the computer can understand.

When the engine sensors detect a warning condition the computer activates the voice synthesiser, which then turns the computer's digital output into the sound of a human voice through a loudspeaker.

A Smooth Ride

Another way in which computers can improve

cars is by controlling the suspension. Lotus Cars is working on a process called 'active suspension'. This technique uses a computer to adjust the stiffness and flexibility of the shock absorbers many times a second so that the car will hold the road well whether it is loaded with passengers and luggage or whether there is just one driver.

Usually sports cars have a stiff suspension so they can hold the road well. But the disadvantage is that the ride is hard and passengers feel every bump. But if a car has a soft suspension to give a 'cottonwool' ride it will not corner as well because the car body can roll too much. Computer-controlled suspension will give the best of both worlds.

Navigation computers are also being developed by Honda and Toyota. These computers tell the driver the direction in which he should be travelling. It can do this by measuring the car's speed, direction, and distance travelled and comparing this with a map of the correct route held in its memory. The driver can then decide whether to turn left or right or go straight ahead according to a series of indicators on the dashboard.

One of the more common computers available in cars measures fuel consumption and can work out the estimated time of arrival. The driver knows how much fuel he is using at any given moment as well as being told by the computer what the average speed of the journey was. Some of the more advanced in-car computers enable the driver to program a cruising speed. The car will then maintain this speed without the driver needing to touch the accelerator.

Computers installed in long-distance juggernauts have a more serious purpose. They serve as an electronic log and enable traffic authorities to determine how long the driver was at the wheel and what his speed was, as well as the distance travelled.

One of the most valuable benefits offered by in-car computers is improved fuel efficiency. BMW, the German car manufacturer, already has a range of cars with a system that determines the optimum mixture of petrol and air required for any given driving condition.

In effect this is like 'tuning' the car so many times a second to give the best possible fuel consumption. The system works by continually measuring the fuel and air mixture and making adjustments to take account of the car's speed, the gear being used, and the engine temperature.

The Future

So what can we expect to see in the future in the way of in-car computers? Theoretically it would be possible for the computer to take over the entire driving function. All the driver would have to do would be to program the car computer with the destination. The computer could then automatically drive the car using information fed to it from sensors buried in the road or by communicating with central traffic computers. Another development we could expect to see is a computer radar that would automatically adjust the car's speed if the vehicle was driven too close to the car in front.

Over the next few years the conventional dial instruments will be phased out. They will be replaced by a monitor screen-style of dashboard with computer graphics. The car driver will be able to call up the electronic display of the information he wants such as engine temperature or fuel level. Driving information could be projected on to the windscreen so the driver would never have to take his eyes from the road. It is also likely that cars will be fitted with computers that would immediately tell the mechanic of any problem. A mechanic would connect the car computer to the garage's pre-programmed computer which would instantly analyse the condition of all the components and identify those that were wearing out or faulty. The photograph below shows the prototype of Honda's dashboard-mounted computerised navigational system





Message Understood

Press a key, and hidden layers of software spring into action – decoding your instructions, searching the memory and all the while scanning the keyboard for your next command

A computer is an assembly of metal, plastic and silicon which, without a program in its memory, is incapable of performing any sort of useful task — just like a record player with no record on the turntable. The process of getting the computer to perform the specific task you require is therefore known as 'programming'. Even the relative beginner to programming will be able to identify two distinct phases to solve a problem. First the problem must be translated and written down in a form which the computer can understand. Secondly, this program must be fed into the computer and 'run'. These two phases can be further subdivided into two stages — the programmer himself goes through the first stage, while in stage two the computer must take the actions (usually without the user's knowledge or intervention).

Suppose that you want to write a program to prepare a payroll. The first thing you need is a perfectly clear understanding of the problem. What do you require as your output from the computer? What information will the computer need to produce the weekly payslips? This may entail salaries, hours worked this week, etc. The next essential is to specify the process by which this output is going to be produced, for example: 'How are tax and pension contributions calculated?'

In a large business application, this may well be done by a trained 'Systems Analyst' — whose speciality is analysing the way a business operates and writing it down in a form which can easily be translated into a program. For home or educational programs, all this would usually be

done by the programmer himself.

If computers could understand plain English then this 'program specification' could be run straight away; but unfortunately they cannot, yet. Many beginners experience difficulties because they try to write the program from start to finish — in the same way that you might translate an English essay into French. Experienced programmers, however, break this stage down even further. They might divide the payroll specification into four 'modules': for inputting the week's data, calculation, storing the cumulative results like 'tax paid this year', and printing out the payslips.

Each module can then be broken down into smaller structures. This is called 'structured programming'; each of these smaller sections is simple and can be expressed as one or two lines in a program. Finally, the whole collection of lines — the program listing — is typed into the computer.

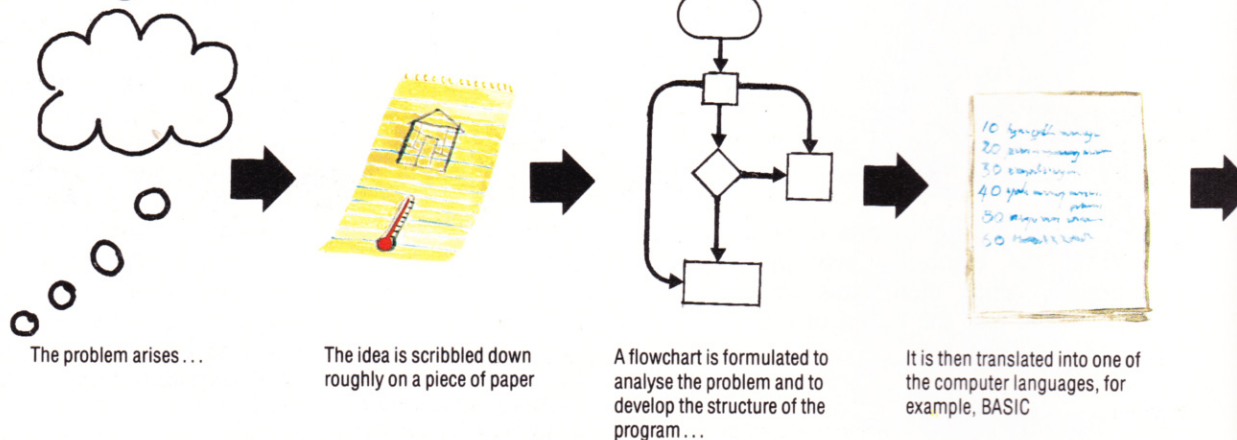
A good programmer always keeps notes from every stage and these reflect the many distinct levels from a problem written in English to a program written in a high-level language such as BASIC.

What happens from the moment you type RUN is entirely under the control of the computer and again involves many different layers or stages. However, the internal operations of the computer are 'hidden' — all the user is aware of is his program asking him for relevant information and producing the required output.

Because the microprocessor cannot understand a high-level language, the prime task ahead of the

From Problem To Program

The origin of a computer program begins with realising there is a problem that needs to be solved; in this case how to keep the temperature of the greenhouse at a constant level. In order to obtain the answer, this problem needs to progress through several layers of processing which results in the completed program



computer is to translate the instructions into machine code. On home computers this is done by the interpreter which is stored permanently in the computer's ROM.

The interpreter is a sophisticated machine code program, executed directly by the microprocessor. When RUN is typed the interpreter starts examining the user's program, character by character. It looks up all the phrases it finds against its own dictionary. If it comes across a character which it doesn't understand (which may simply be because you made a typing mistake) it will stop trying to interpret the program and print a message on the screen such as SYNTAX ERROR.

If the word is in the interpreter's dictionary (e.g. PRINT) this immediately passes to the part of the interpreter which knows how to deal with that function. In this case the routine will now examine what comes after the word PRINT in the user's program and prepare this data as a stream of characters to display.

This is where the next level comes into operation. Somewhere else in the computer's memory is a routine which can accept a stream of characters, store them in another area of memory reserved for the screen, and arrange for them to be converted into the kind of signals needed by the television screen or monitor. This is something that has to be done continuously even while the program itself is purely engaged in calculation.

The same is true at the other end of the computer — the keyboard. A specially written program routine within the computer has to scan the keyboard to find out whether any keys have been pressed, and if so place the appropriate codes in another area of memory for use as the input to the user's program. And because you might want to halt the operation of the program at any time using the BREAK key, the keyboard must be scanned continuously — even while the program is running.

In fact the microprocessor found in most home computers can only do one job at a time so it effectively has to share its time between interpreting the user's program and its own




internal functions, such as checking the keyboard and controlling the screen. One method is 'interrupt-driven' where a special electronic circuit interrupts the microprocessor perhaps 50 times every second and 'reminds' it to perform its housekeeping tasks and other functions on the screen and keyboard, before resuming what it was doing.

So even when your program has been typed in, many levels of processing have to be carried out by the computer before the results are produced. Though the process may seem a complex one most of it is looked after by the computer.

Nowadays the trend is moving towards user-friendliness, with the computer doing as much of the routine work as possible. The next generation of computers will be able to write the whole program themselves from a plain English specification.

The Hidden Software

In any computer, there is a complex hidden software hierarchy continually working. Among its many tasks, it monitors and checks when and which key is being pressed, what is on the screen, what instructions are being given to the peripheral, and the status and content of the RAM memory. All these functions are continually in progress while the operator is simply concerned with the next function of his program. The principle of the hidden software hierarchy remains the same, in both sophisticated business computers (as pictured here) and in reasonably-priced home computers



The program is then entered into the computer's memory via the keyboard

The BASIC program then passes through an array of chips. These convert it into machine code that is relayed to the CPU

The computation is then made in the CPU. The resulting data are subsequently transmitted to a peripheral (e.g. a printer, monitor, disk drive etc. . .)

In this case the printer produces a printout or 'hard copy'

If the program is accurate, the problem is now solved

True Or False?

Computers may not be able to 'think' yet, but they can certainly follow the laws of logic

The CPU (Central Processing Unit) is often described as the heart of a computer. It's the place where all the computations and logical decision-making take place. But how are these decisions and calculations made?

To understand, we need to know the basics of binary arithmetic and be familiar with logic gates. In computers, these gates are simple electrical circuits able to make logical decisions and comparisons. This may sound more complicated than it is, and the principles can easily be illustrated using examples from everyday life.

There are three fundamental types of gate — the AND gate, the OR gate and the NOT gate. Capital letters are used when writing about gates to differentiate the words from the usual English 'and', 'or' and 'not'.

Logical Connections


An AND gate is a circuit that gives a 'true' output if all the inputs are 'true'. Let's see what this means. Suppose you would like a trip to the country. If you have a car AND some petrol, you can have the trip. If you have some petrol but no car, you can't go. Similarly, a car but no petrol means no trip.

In this AND 'circuit', there are two input conditions and both need to be 'true'. To get the trip (the 'output') it must be true that you have a car AND it must be true that you have some petrol. Then the output becomes 'true' — it is true that you get the trip to the country. Later, we'll see how this logic diagram can be shown as a logic equation and also how it can be represented in a 'truth table'.

Imagine a slightly different situation. Somebody would like to take a trip to the country. The trip will be possible if they have either a car OR a bicycle (we'll assume the car has petrol this time). If they have a car, they can go. If they have a bike, they can go. It is only if neither of the input conditions is true that the trip becomes impossible — in computer jargon, the output becomes false (i.e. it is not true that he gets the trip to the country).

There is one more essential logic gate to consider, the NOT gate. This gate simply gives as an output the opposite of the input. If the input is true, the output will be false. If the input is false, the output will be true. Extending our metaphor of a trip to the country, whether by car or by bike, it must be false that there is a flat tyre to get a trip to the country. If the input (a flat tyre) is true, then the output (a trip) will be false.

AND




A trip to the country will be possible (true) if there is a car AND petrol

OR



A bicycle OR a car allows a trip to the country to be made


AND/OR



AND and OR logic elements can be combined for more complex situations. Computers incorporate thousands of such gates

NOT

A NOT gate gives an output that is the opposite of the input. If it is true that the tyre is flat, it is not true that there will be a trip to the country



ANDY LESLIE/MARK WATKINSON

These logic elements can be combined and we have illustrated this using the trip to the country example. Combinations of AND, OR and NOT allow all decisions based on conventional logic to be made. It is interesting to work out the logic decisions (gates) needed for other problems. Try working out what would be required for a garden barbecue, for example. It can become quite complicated. To have the barbecue in the garden (the true output) we would need various input conditions: money OR a cheque book OR a credit card (to buy the food and drink) AND a free evening AND fine weather AND a grill AND charcoal.

In computers we use the binary digits zero and one to stand for false and true respectively. The computer interprets a plus voltage as one and a zero voltage as a zero. An AND circuit can easily be made using transistors so that if both inputs are plus voltages, the output will also be a plus voltage. If either or both of the inputs is a voltage of zero, the output from the circuit will also be a zero.

An electronic OR circuit gives a positive voltage output if either or both of the inputs is positive. If both the inputs are zero, the output will also be zero. In a NOT circuit, the input is simply reversed: if the input is positive, the output will be zero; if the input is zero, the output will be positive.

Truth Tables

The symbols we have used in the illustrations are the same as those used in computer circuit diagrams. To see just how easily logical decisions can be implemented using electrical circuits, let's look at the 'truth table' for the AND illustration. If we use the letter c to represent the 'having a car' input condition and the letter p for the 'having petrol' input condition, we can represent the 'having a trip to the country' output condition using a t. We can then use T to stand for true and F to stand for false. The truth table shows all possible combinations of input conditions and the effect of using AND on the output. It looks like this:

CAR	F	T	F	T	(c)
PETROL	F	F	T	T	(p)
TRIP	F	F	F	T	(t)


The truth table for 2-input AND

CAR	0	1	0	1	(c)
PETROL	0	0	1	1	(p)
TRIP	0	0	0	1	(t)

The same truth table using 0 and 1 for False and True

The 7408 Chip

Large silicon chips often contain thousands of gates to perform AND, OR and NOT logic. The small 7408 chip illustrated contains all the transistors and circuitry for four AND gates. The gates are represented using the logic symbols, rather than showing the actual circuits



TONY LODGE



The Lynx

This elegantly styled micro has an exciting range of graphics and one of the best keyboards available

The Lynx is a British-designed computer manufactured in Cambridge by Computers. It is larger and heavier than its two main British rivals, the Sinclair Spectrum and the Oric, but is slightly smaller than the American-made Commodore and Atari machines. The Lynx is one of the smartest-looking computers in its class and is finished in a business-like shade of grey.

It has 48 Kbytes of memory which can be expanded to a massive 192 Kbytes. If you take into account that a professional business computer has usually at least 64K of memory you will appreciate how versatile the Lynx can be. The Lynx has a full size typewriter-style keyboard that gives superb tactile feedback (you instinctively know whether you have pressed the right key).

The Lynx can display eight different colours and has 32 graphics characters stored in its memory, though not marked on the keyboard. It comes with its own version of BASIC, which includes several useful commands, including AUTO which automatically numbers program lines.

The central processing unit of the Lynx is the Z80 Microprocessor which has been used in some very sophisticated microcomputers.

Lynx Keyboard

The Lynx looks workmanlike in its hard plastic case. It has a professional keyboard with 57 keys and a space bar. The alphanumeric keys are arranged in the standard QWERTY fashion. The pairs of arrow keys at both ends of the third row from the top are cursor control keys for use with the screen editor only. The ESCAPE and CONTROL keys are at the top left and the BREAK and DELETE keys at the top right

Parallel Interface

This socket allows peripherals using parallel communications, such as a disk drive, to be connected to the Lynx

Power Socket

This DIN socket connects the DC voltages from the power supply unit to the computer

Loudspeaker

All the sounds that the Lynx can generate are produced by this loudspeaker

Keyboard Socket


A ribbon cable connects the keyboard to the main computer circuit board here so that the computer can determine which keys are pressed

RAM

These chips provide the random access memory that is available to the user for storing programs and data

Microprocessor

The Lynx's central processing unit is a Zilog Z80A



John Shirreff

The brains behind the Lynx range of computers belong to 35-year old John Shirreff. Not long after leaving Cambridge University in the late 60s, he designed and built inflatable portable structures, and played drums with a number of rock bands. He continued work as a sound recordist and as a session musician while developing his interest in electronics. The two interests came together when he met Dave Vorhaus who ran a recording studio. In 1981, he designed a new Z80-based 64 Kbyte expandable business computer. The design team's experience with the Z80 processor on this project provoked much of the thinking behind the Lynx range

**Cassette Socket**

This QIN socket is where the cassette unit is connected to the computer

RS232 Port

This socket allows peripherals using serial communications, such as a modem or acoustic coupler, to be connected to the Lynx

RGB Interface

A colour monitor can be connected to the Lynx using this socket

Modulator

The signal from the video chip is converted by this to a form that can be accepted by a television set

Clock

The electronic 'beat' of this oscillator is used to time and synchronise all the computer's operations

Video Chip

This chip generates the video signal that can be fed directly to the RGB interface for display on a colour monitor and to the modulator when a television is used as the display unit

Input/Output Chip

This chip converts inputs to the computer to its internal form and its outputs from this internal form to a form suitable for the receiving device

BASIC ROMs

The Lynx's BASIC language is stored permanently in this pair of ROMs

LYNX**PRICE**

£225

SIZE

350mm x 213mm x 60mm

WEIGHT

1564g

CPU

Z80A

CLOCK SPEED

4 MHz

MEMORY

48 Kbytes RAM expandable to 192 Kbytes
16 Kbytes ROM providing BASIC and monitor

VIDEO DISPLAY

Text mode of 24 rows each of 40 characters
High resolution mode with 248 x 256 dots

INTERFACES

Television connector, RGB video, Composite video, RS232 cassette connector, parallel expansion socket

LANGUAGE SUPPLIED

BASIC

COMES WITH

Aerial lead, cassette lead, power supply unit, manual, cassette

KEYBOARD

Professional, QWERTY keyboard with 57 keys and space bar

DOCUMENTATION

Although the manual provides an acceptable introduction to BASIC it lacks the kind of information that the more advanced user would want. The explanations of the Lynx's sound and graphics facilities are adequate for the beginner. But the introduction to machine code and the use of the monitor is far too sparse to have any value to the newcomer.

Unconnected topics are often grouped together. This, coupled with the lack of an index makes it difficult to locate some topics.

The illustrations are reproduced directly from sketches by the author. This is acceptable for the cartoons, but is unsatisfactory for the technical illustrations



Consulting The Chip

Specialists will soon be freed from routine by 'Expert Systems': computers programmed to analyse complex data and answer questions about it

Artificial intelligence, the creation of computers that think and make decisions like their human creators, is still a science-fiction fantasy. Full understanding of the human brain and its working is a mammoth task and, although some advances are being made there is little prospect of a '2001-style' intelligent computer for many years.

highly and expensively-trained professional scientist. Thanks to the computers these valuable people can now get on with more original work.

But expert systems have something more to offer than simply replacing human experts. Once the specialist knowledge is being used by the program, the computer often throws up some unexpected facts. Sometimes the machine spots relationships between items of information that the humans have missed, and suggests new avenues that can be explored.

So it is generally agreed that expert systems are, or at least will be, an important development in computer applications. If the program in charge of the system works properly, the computer can act like an expert witness. And many computers can use the same program, turning the specialist knowledge of a single person into a large number of equally expert computers.


The obvious problem for the researchers is writing a program that *does* work properly: one that is as 'clever' as a human specialist.

Creating The Program

The first step is to think about how human experts make decisions about evidence and questions concerning their speciality. Human thinking is not particularly logical, certainly not when compared with the way computers work, and it depends very much on experience. If a new problem or question is put before a human expert, it is mentally compared with the large number of other situations the expert has encountered. Then, by comparing the new situation with those already in the memory, some tentative conclusions can be made and the appropriate action taken.

But representing the immensely detailed knowledge of an expert like, say, a doctor in these structures means that an enormous number of rules need to be stored and linked in very complex ways. And further modifications are also necessary if the computer is to mimic human behaviour. Human doctors are rarely certain about anything, and can only say that they are 'almost sure' or 'pretty confident' about an opinion. Based on just a couple of symptoms, our doctor might only be 30 per cent sure of a diagnosis.

So the rules in the computer model have to have probability values attached to their conclusions, ranging from 100 per cent where there is only one conclusion possible to 1 per cent



COURTESY OF STANFORDS LTD.

A View From Space

Landsat 4 was launched in July 1982 and is in an orbit that covers the whole surface of the world returning over the same spot every 20 days. All the data from the sensor is sent in digital form and by using computer techniques objects only 40 metres across can be resolved and geographical features interpreted. In the illustration the digital information has been photographically processed to show aspects of London and the southeast of England. Clear water is dark blue, shallow water with sediments is light blue, towns and ploughed fields are blue-grey, heathland red-brown, corn at harvest green and other vegetation is bright red.

But if the task can be limited, if a computer only needs to seem 'intelligent' in a very restricted field of human activity, then reproducing at least the appearance of intelligence becomes much easier.

This is the theory behind expert systems. The idea is that an expert in a particular field, like a top geologist or surgeon, can feed expert knowledge and rules for dealing with it into a computer system. Then the computer program handling the knowledge and the rules is open to enquiry from unskilled people, who can type in questions about its specialist field and receive meaningful answers.

Expert systems could be useful in many ways. One program has already been developed to diagnose the cause of stomach pains by questioning patients about their symptoms. Another uses our knowledge of geology to pinpoint the most probable sites for finding molybdenum or other minerals. And a third deduces likely structures for organic molecules from masses of unstructured experimental data. All these tasks would normally be taken on by a

when there are 100 equally likely conclusions.

Most expert systems work by starting a dialogue between questioner and computer. The questioner has to enter details of the problem that the computer expert is to solve, and the easiest way of avoiding trouble here is to make the computer ask the user a series of multiple-choice questions. This avoids the problem of the user entering words or sentences that the computer does not understand. Then the computer tests the entered information against the rules in its store of knowledge. The route the computer takes through

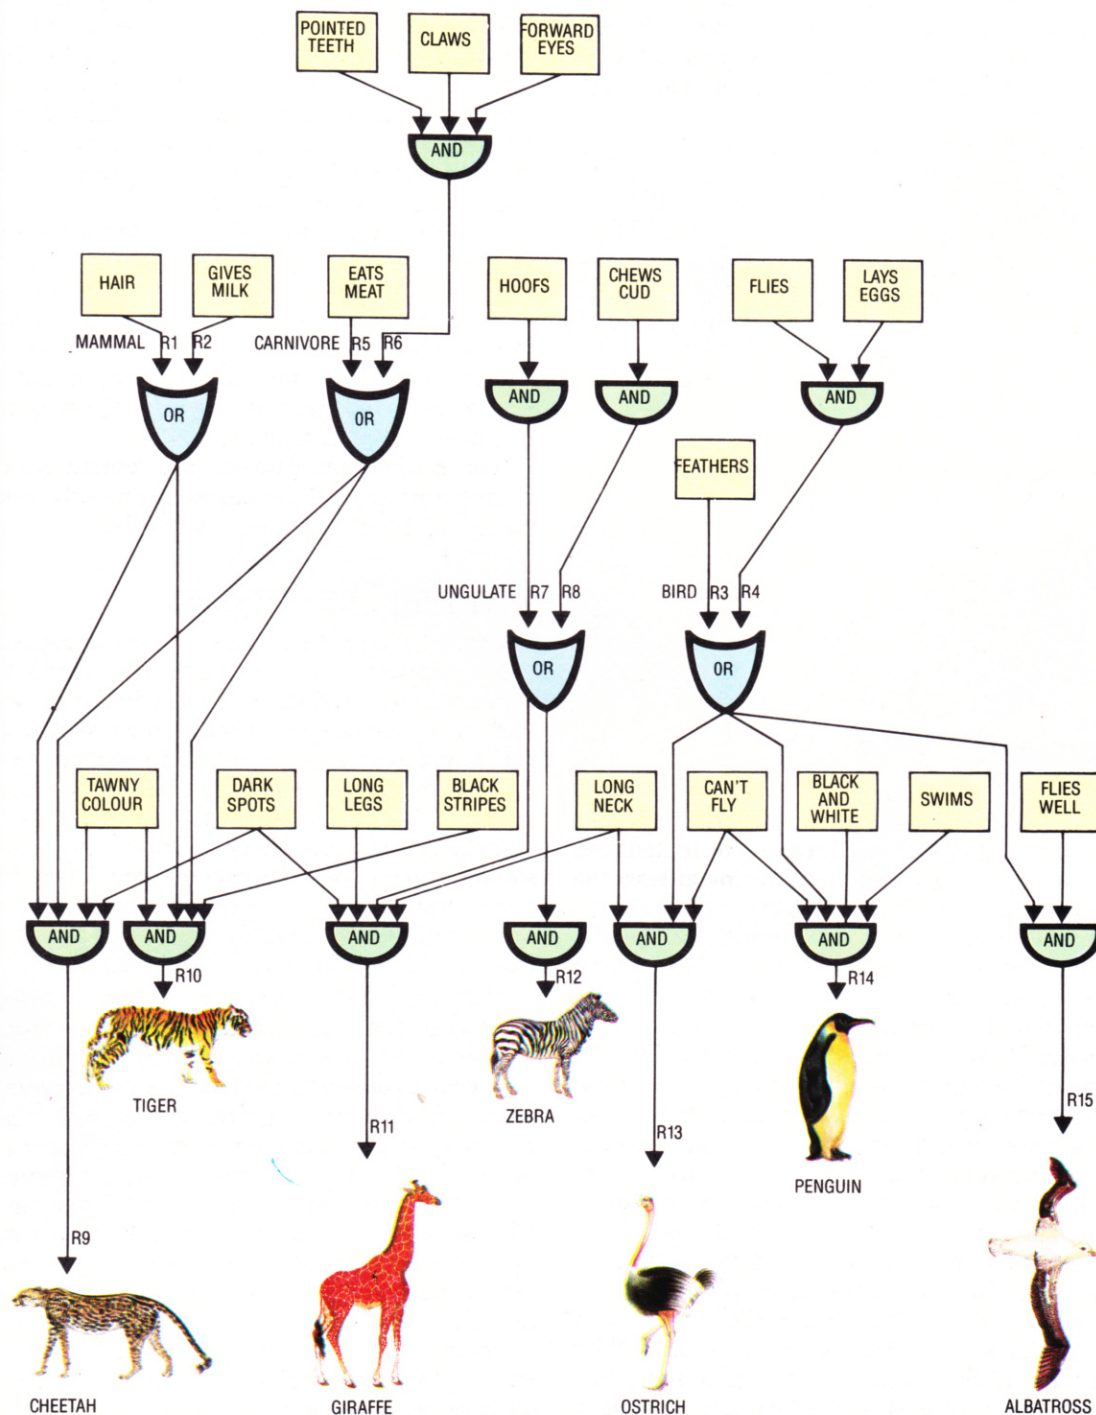
the rule network depends on the answers the user gives to its questions, and as the program progresses through the network each step determines the computer's next question.

The best-known expert system in Britain, the Mickie medical diagnosis program, runs on a microcomputer. And other commercial products are on the way that will turn the home machine in the living room into an expert on several subjects.

We might not be able to talk to our computers yet. But we can now ask them questions and receive answers we have reason to trust.

A Network Of Animals

This is a chart showing how the 'knowledge' of an expert system is put together as a series of simple rules. Here the rules are numbered from R1 to R15, and the purpose of the rules is to identify an animal from any information you have about it. For example, suppose you enter the information that the animal has pointed teeth. The computer starts at the box at the top left-hand side of the chart and sees where it can go in the network, aiming towards one of the animals named at the bottom. The AND in the circle means that all three of the conditions in the boxes leading to it must be true before R6 can be applied. The computer is unable to progress further until it is given further information and needs to ask two questions: does the animal have claws? and does it have eyes in the front of its head (forward eyes)? If the answer to both is yes, then R6 allows the computer to move down the chart. But R6 leads to an OR, which means that it can only move on if R6 is true or R5 is true. Here R6 is true, so the computer goes down further. Here there are two AND possibilities, each leading to different conclusions. For either R9 or R10 to apply, and give the answer CHEETAH or TIGER, the computer has to ask more questions. In this case it has to ask whether the animal has black stripes or dark spots. With this information it can then provide the answer. This type of chart is typical of the way in which the 'Animals' computer game works





Fit To Print


Daisy wheels, ink jet and dot matrix printers are among the newest developments in print technology that are finding their way into the home and office

The Daisy Wheel

The daisy wheel has the characters attached to the end of the 'petals'. Hitting a key moves the wheel to the appropriate petal so the character can be printed

Paper Guides

The paper is fed through the printer by a rotating spiked wheel that catches the perforations on the edge of the paper and feeds it through the printer



Ribbon Cartridge

Most daisy wheel printers have a ribbon cartridge that can be changed in a matter of seconds. The new ribbon is enclosed in a plastic case that slots into the printer

The Mallet

A small metal 'mallet' hits the character at the end of the daisy wheels and pushes it against the ribbon, causing the print to appear on the paper

Ribbon Tensioner


Turning this control will tighten or loosen the ribbon. This is useful when changing ribbons as the new ribbon will often become slack in its cartridge before use

Control Panel

This panel has two main controls. These turn the printer 'on line' which prepares it to receive data from the computer. It also moves the paper to the top of the next page if the printer stops in the middle of a page. Other controls are a 'line feed' which moves the paper up a line at a time

Daisy Wheel Motor

The motor revolves the daisy wheel and aligns the characters with the mallet which strikes them against the ribbon and paper




The Dot Matrix Printer

The dot matrix method uses a grid (matrix) of dots to form a character. The print head contains a group of needles that punch against the printer ribbon leaving a dot on the paper. As the print head moves along the paper, the pins are fired by electrical impulses in the correct sequence to form the character

The Ink-Jet Printer

This appropriately-named printer fires a jet of ink through a nozzle which breaks the ink up into tiny drops. Each drop is given an electrical charge and passes through metal deflector plates. The charge ensures that the ink drops hit the paper in the right pattern to form the character



Up to now you probably haven't given much thought to using a printer. After all, if you are quite happy using your home computer to play games or calculate your home finances there isn't much need for a printed copy of what your television screen or monitor displays.

But as you become more skilled at using a home computer the limitations of doing serious work without a printer become obvious. If you are interested in writing your own programs, you will want to keep copies of your program lists. If you use your computer for your accounts, a printed record of the calculations will be needed.

Choosing a printer for your needs is a fairly tricky job. How much you have to pay depends on the speed at which the printer can produce words and the quality of the results.

Choosing Your Printer

There are three main types of printers for home computers: the dot matrix, daisy wheel and thermal printer.

The most common method of printing is the dot matrix. This works through a print head that contains a group of needles. Characters are printed by combinations of these needles striking the ribbon. The advantage of the dot matrix method is that it is very fast and the printers are relatively inexpensive at around £300. However, because the letter or number is made up of a series of dots, the print quality tends to be poor. The printer is also rather noisy.

Some dot matrix printers overcome the problem of poor print quality by overprinting the dots two or three times. In this case, the print head moves slightly so that the new dots fill in the spaces between the dots that were first printed.

Dot matrix printers are acceptable if you want to keep rough copies of what your computer has produced. They will also produce charts and graphs, as the print head can be made to print patterns as well as characters. But you will need a different printer if you need high quality printing for, say, a letter to your bank manager.

For typewriter-style quality you must turn to a daisy wheel printer; so-called because the printer uses a wheel with long 'petals' that look like a daisy. At the end of each petal is a letter, symbol or number. To print out, the wheel turns to put each petal in line with a little metal 'mallet' that pushes the character at the end of the petal against the printer ribbon and paper. The daisy wheel is made of either plastic or metal. You may also change daisy wheels for different varieties of print face as you would with a golfball typewriter.


The problem with daisy wheel printers is that they are much slower in operation than dot matrix printers and tend to be more expensive. They are also not as suitable for charts and graphs, as several wheels would be required to produce the various graphics shapes. You can expect to pay at least £400 for a daisy wheel.

Higher up the price scale are ink jet printers.



The 'Ball-Pen' Printer

This represents a new development in printing techniques. The print head holds four specially designed ball point pens. When the PRINT command is given, the paper moves up and down to create the vertical strokes in the character while the pens move sideways to create the horizontal strokes. The advantage of such a system is that it can be used for printing coloured charts and graphs. It also offers higher print quality than the dot matrix method as its characters are formed by single pen strokes. However this system is comparatively slow and the pens need to be replaced regularly if long pieces of text are continually printed



STEVE CROSS

These eject drops of ink that form the shape of the character to be printed. The ink is forced through a nozzle which breaks it up into tiny drops. These drops then pass through an electrode and are given electrical charges. A pair of metal plates then deflects the drops in different directions, so as to form the shape of the character. Ink jet printers are so fast that they can print around 20 metres of characters every second!

An alternative method, the thermal printer, uses heat-sensitive paper. The print head transfers its heat to the paper in such a way that the paper turns black in the area touched to form the appropriate character. Thermal printers are very quiet and fairly fast. One of the most popular is Apple's appropriately-named Silent Type model. Thermal printers are reasonably priced at approximately £300, but you will have to use special heat-sensitive paper that is more expensive than ordinary paper, and the print quality is not as good as that from the daisy wheel.

The Right Interface

When you decide to buy a printer, as well as determining exactly what you want to use it for, in terms of quality and speed, you must also be certain that it can be used with your particular computer. The connecting plug from the printer must be compatible with the computer. The socket for plugging in a printer is usually found on the back of the computer and is called an 'interface'.

The three most common types of interface are Centronics, IEEE488 and RS232. Centronics is also referred to as a 'parallel' interface. Your computer will have an opening carrying at least one of these three interface names.

However, the computer industry is notorious for its incompatibility, and you may find that a printer and a computer using the same connection might still not be compatible. This is because the interface must be set to work at the same speed on both the micro and the printer. This speed is known as the 'baud' rate and is the speed at which bits from the computer's memory can be transferred to the printer. The bits are sent to the printer in one of two ways: either they follow each other down a single wire, as in the RS232 interface ('serial' method) or they are transmitted together down several wires, as in Centronics and IEEE488 interfaces ('parallel' method).

Printers use two main methods of handling paper. Paper can be fed one sheet at a time as with a typewriter. Alternatively it can be 'tractor'-fed with the two sprockets catching the paper in perforated holes on each side of the sheet; in much the same way as a camera moves the film from frame to frame. Tractor or sprocket-fed paper is more convenient as you can leave the printer to feed itself with paper. However, this method will not accept headed paper which must be fed into the printer one sheet at a time.

When buying a printer, you should decide what you wish to use it for and then choose the best model that you can afford for your computer.

Routine Matters

Programs within programs: we introduce a new feature of Basic that will keep your programs neat and manageable

In previous instalments of our Basic Programming course we have typed in programs, run them, made modifications to them and then cleared the memory (using the NEW command) when we wanted to enter new programs. When we have needed to run the old program again, it has been necessary to type the whole thing in again.

To save this repetitious work, all versions of BASIC are provided with a command that allows any program to be stored on cassette tape. The program below can be saved on tape by using the simple command SAVE followed by a file name. The program calculates the number of tiles needed to tile a room.

```

10 REM THIS PROGRAM CALCULATES THE
  NUMBER OF TILES
20 REM NEEDED TO TILE A ROOM
30 PRINT "ENTER THE SIZE OF SIDE OF TILE IN
  MM"
40 INPUT A1
50 REM LINE 60 FINDS AREA OF TILE
60 LET A2 = A1 * A1
70 PRINT "ENTER THE NUMBER OF WALLS"
80 REM W SETS LIMIT FOR LOOP
90 INPUT W
100 FOR X = 1 TO W
110 PRINT "LENGTH OF WALL NO. "; X; "IN METRES"
120 REM D IS DIMENSION OF WALL
130 INPUT D
140 REM IT IS CONVERTED TO MM
150 REM IN THE SUBROUTINE
160 GOSUB 380
170 REM LINE 190 SETS L TO
180 REM LENGTH OF WALL IN MM
190 LET L = D2
200 PRINT "HEIGHT OF WALL NO. "; X; "IN METRES"
210 REM LINES 230 TO 250 SET H
220 REM TO HEIGHT OF WALL IN MM
230 INPUT D
240 GOSUB 380
250 LET H = D2
260 REM LINE 270 SETS A3 TO AREA OF WALL
270 LET A3 = L * H
280 REM S (SUB-TOTAL) IS AREA OF WALL DIVIDED
290 REM BY AREA OF TILE
300 LET S = A3/A2
310 REM T (TOTAL) HAS THE NEW SUB-TOTAL
320 REM ADDED EACH TIME THRU THE LOOP
330 LET T = T + S
340 NEXT X
350 REM PRINT THE TOTAL
360 PRINT T

```

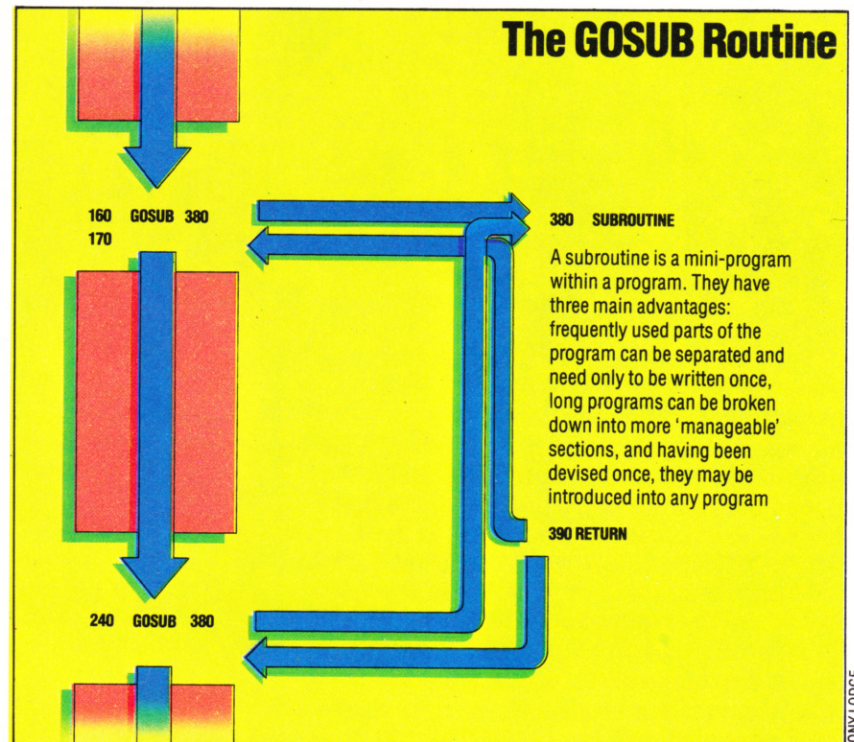
```

370 END
380 LET D2 = D * 1000
390 RETURN

```

Having typed in the program, all you need to do to save it on cassette tape is to use the SAVE command. First, of course, the cassette recorder must be set up in accordance with the instructions in your computer's handbook. The SAVE command is extremely easy to use. Just type SAVE followed by a file name in double quotation marks. A file name is the name given to a file, and a file, in computer terms, is like a file in a filing cabinet — a program or set of data that can be stored away or retrieved when required. It is best to use a file name that will remind you of the function of the program. Since our program calculates the number of tiles needed to tile a room, we could call it "TILES". Once the cassette recorder is set up, insert a blank tape to hold the program.

Cassette decks with a remote control socket can usually have the motor controlled directly by the computer. Otherwise, set the recorder in the record mode and then put it in the pause mode. Type in the SAVE command, including the file name. Set the recorder running by releasing the pause control and then hit RETURN.



To test that the program has been recorded correctly, erase the computer's memory by typing NEW<CR>. Rewind the cassette, put it into the play mode and then load the program back into the computer using the LOAD command. LOAD must be followed by the file name of the file wanted. Type LOAD "TILES" and then hit RETURN.

After the program has been loaded into the computer, a message on the screen such as READY or OK indicates that the load has been completed. LIST the program and check that it is the same as the one you typed in.

GOSUB

GOSUB is a statement that diverts the flow of a program to a subroutine. A subroutine is like a separate mini-program or program-within-a-program. In the program used to illustrate it here, the subroutine is very simple. It is included to show the principle, although other ways could easily have been devised to produce the same results without using a subroutine.

Our program calculates the number of tiles needed to tile a room by finding out the area of the tiles used. It then asks for the length and height of each wall to be entered. It works out the area of the wall after the length and height has been converted from metres to millimetres. The number of tiles needed is found by dividing each wall's area by the area of a tile and adding the results. The conversion of wall length and height into millimetres is done in the subroutine, which simply multiplies the length or height (in metres) by 1000 to find the equivalent in millimetres.

Subroutines have three advantages. Frequently used parts of programs can be separated off and only need to be written once — no matter how often the operation is required. They allow long and complex programs to be broken down into more manageable and easily understood units or sections. Finally, subroutines can be re-used in any program where its function is appropriate.

In our program, the subroutine starts at line 380 and consists of only one statement: LET D2 = D * 1000. This takes D, the wall dimension (length or height) and multiplies it by 1000 to convert from metres to millimetres. The result is assigned to variable D2.

The instruction that forces the program to go to the subroutine is GOSUB. It occurs first in line 160. Variable D was assigned the value of the length of the wall in line 130. Line 160 forces the program to go to the subroutine, where variable D2 is given the value of D multiplied by 1000. The RETURN instruction in line 390 is needed to make the program return from the subroutine to the main program. Subroutines always return to the line after the GOSUB statement, in this case, to line 170.

The next occurrence of GOSUB is in line 240, which 'calls' the same subroutine. This time, the subroutine RETURNS to line 250. Although this program uses only one subroutine, it is possible to use as many as are needed. In every case, the GOSUB statement will have to include the line

number of the appropriate subroutine. Notice that the END statement occurs in line 370, before the subroutine. END indicates the end of the main program and also serves to stop the program from running on through the subroutines after it has been completed.

Although this program is a little longer than previous programs in this course, it is really no more complex. Try and follow it through, line by line, and see what is happening at each stage. Apart from GOSUB and subroutines this program introduces only one new concept — longer variable names.

It may be helpful to draw boxes with the variable names written on them and to write in the values at each stage.

Line 300: LET S = A3/A2 will sometimes give a number with a decimal fraction. Try running the program and entering the tile size as 110mm and the wall length and height as 2.3 and 1.8 metres respectively, using just one wall. You should get an answer of 342.149 tiles. Since tiles are never sold in units of less than one, this answer is not completely appropriate. Next time we will look at one of the ways of getting an appropriate answer in whole numbers.

Exercises

- See what happens if you enter the size of the tile as 0mm. You should get an error message at the end of the run. Why is this? Why don't you get a similar error message when you enter the length of one of the walls as 0 metres? Hint: multiplying by zero and dividing by zero are not the same thing — try it on your calculator!
- The program only works for square tiles. See if you can change lines 30 to 60 to find the area of rectangular tiles (just as we found the area of rectangular walls later in the program).
- Add a statement at line 355 to increase the total number of tiles by five per cent to allow for wastage. Multiplying a number by 105/100 will increase it by five per cent.

Basic Flavours



This is not available on the Spectrum, which means that some other method is needed to bypass lines 380 and 390. Changing line 370 to: 370 GOTO 400, and adding 400 PRINT "END" will achieve this




The Spectrum requires all variables to be defined before an arithmetic operation can be performed. So that in order to make line 330 work properly, a new line must be added and should read:
5 LET T=0



Appears as two words on the Spectrum, although only one keystroke is required

When Minus Equals Plus

Computers like their circuits simple and so must employ a clever trick to perform subtraction by means of addition



Voyager 2

The spectacular Voyager was the first space explorer to go beyond our solar system. It travelled through space, taking pictures and gathering information, while the on-board computer converted the data into binary digits. The data were sent back to earth at the staggering rate of 116,000 bits per second. They were then processed by the NASA computer in Houston, Texas

In the first part of this series, we discovered that binary digits could be used to represent any decimal number. Binary numbers have the disadvantage of being longer than their decimal equivalent, but are convenient for the computer as the zeros and ones may be represented by zero and positive voltages. We also saw that binary numbers can be added together very simply.

On paper, binary numbers can be subtracted as easily as decimal numbers, following the same rules used in decimal subtraction. Computer designers realised long ago, however, that adding circuits (electronic circuits that perform addition)

could both add and subtract without the need for special subtraction circuits. We shall find out how this is done.

Two's Complement

One method of representing negative numbers in computers is known as 'Two's Complement'. With this the process of subtraction appears as just another part of the addition. Consider the following arithmetical problem:

$$\begin{aligned} 16 - 12 &= 4 \\ \text{or } 16 + (-12) &= 4 \end{aligned}$$

Here 12 is taken away from 16 but the process of subtraction can equally be seen as an addition: the addition of 16 and negative 12. In both cases the answer is the same and the only difference is the use of arithmetic signs and brackets. This slight modification can be used by the computer to both represent negative numbers and simplify the problem of subtraction.

For simplicity assume our computer is only large enough to handle 5 digits. Of course real computers can handle numbers with thousands of digits. Our 5-digit computer adopts a method of working: the leading digit on the left hand side is to be considered separately from the other 4. If the leading digit is 1 it is to represent negative 16 and if it is 0 then it is of course a zero. The remaining 4 digits are positive and follow the binary conventions we saw in the last instalment of the course

[]	[]	[]	[]	[]
-16 or 0	8s	4s	2s	1 or 0

So for example the binary number 01000 is decimal 8 and 10000 is decimal -16. But what about 10100? This includes -16 and +4 giving -12.

How many numbers can be represented with only 5 digits using this convention? The largest positive number is 01111 or decimal 15 and the greatest negative number is 10000 or -16. With a little experimentation you will see that every number between -16 and +15 can be represented.

Binary Decimal

10000	-16
10001	-15
10010	-14
10011	-13
10100	-12
etc.	
11111	-1
00000	0
00001	1
00010	2
etc.	
01110	14
01111	15

If we increased the number of digits our computer was able to handle, we could of course expand the range of numbers.

Early on in the development of binary computer arithmetic, a very simple trick was discovered for finding the Two's Complement, or negative form, of a number. There are two steps to this trick.

First, invert each digit. So whenever you see a 1 put a 0 and wherever there is a 0 change it to a 1. Secondly, add 1 to the reversed number.

Follow the method as it is laid out in the example below. We are using +12, the binary equivalent being 01100. (The leading 0 on the left hand side is not strictly necessary as 01100 is the same as 1100. But since our computer has 5 digits we must fill up every one).

01100 (= +12)

First Step: 10011

Second Step: 00001 (+1)

10100 (= -12)

Now let's look at how our computer tackles a problem of subtraction; for example: 12 minus 4.

+ 12 is	01100
- 4 is	11100 (using Two's Complement)
12+ (-4)	101000

Notice that we now have 6 digits. Since our computer is only large enough to register 5 digits, the leading left hand digit is called an overflow digit and is ignored, leaving 01000 or 8 in decimal, the correct answer! A slightly more complex example is: 4 minus 12.

+4 is	00100
-12 is	10100
4+(-12)	11000

As a final example let's try dealing with two negative numbers together: -3 -4 = -3+ (-4) = -7

3 is	00011
so -3 is	11101 (using Two's Complement)
and -4 is	11100
	111001

Again we are left with a 6-digit number. Once the overflow is discarded, we have the binary number: 11001 or -7 in decimal.

These subtractions used only addition and the trick of the Two's Complement (which itself uses only reversal of digits and addition). The advantage to the computer is that binary digits can easily be reversed using a NOT Gate (see page 68).

A NOT Gate has one input and one output. It is a very 'perverse' Gate because whatever value you feed in, the output is the opposite. So if the input is 0 the output is 1 and if the input is 1 the output is 0. This characteristic of 'inverting' is exactly what is needed for the first (the inversion) step in the trick of the Two's Complement. In the next part of the course we will see how addition can easily be carried out by a computer using a combination of logical Gates.



Dots And Dashes

Morse code is one of the earliest illustrations of binary coding in electronics. In 1837 the first electric telegraph was laid in London with two miles of cable joining Euston to Camden Town railway stations. Later the same year in America, Samuel Morse demonstrated his celebrated code for transmitting messages. Each letter was a combination of two signals: dots and dashes

WIN A B.B.C. 'MICRO B' HOME COMPUTER

Just look at what the 'Micro B' has to offer.

- Recommended by the Government for use in primary and secondary schools.
- Uses the BBC version of the most popular computer language – BASIC.
- Loaded from any standard audio-cassette recorder.
- Has 73 key full typewriter-style keyboard with 10 easy-to-understand function keys.
- 32 K Random Access Memory (RAM) gives high definition graphics.
- 3 part music synthesiser capability.
- Standard in-built connectors for domestic T.V./monitor/audio-cassette/computer accessories.
- Easily adaptable to other computer languages (e.g. Pascal, FORTH and LISP).
- Flexible additions include word processor, disk drive and Micronet/Prestel adaptors.



10 First prizes of BBC 'Micro B' computers **PLUS** **50** Sinclair Spectrums for runners-up

HOW TO ENTER:

Answer the questions below (issue 3 contained a glossary of terms to help you) and complete the tie-break statement. Then send your answers on a postcard or the back of a sealed envelope to Dept. H.C.C. Orbis Publishing Ltd., 20-22 Bedfordbury, London WC2N 4BT. Competition closes 6 weeks after publication of this issue. Please also tell us which computer you already own.

ALL YOU HAVE TO DO:

1. Give the answers to the following definitions

- A.** A group of four bits, or half a byte.
- B.** A collection of bits (usually eight) that make up the smallest amount of usable information in a program.
- C.** An impact printer mechanism with the letters and numbers on spokes attached to a central stem.

- D.** A set of instructions that are frequently required during the running of a program and that can be called up at any point in the program.
- E.** A number system on a base of 16 that is written in both letters and numbers.
- F.** A correction made to a computer program, usually as a group of instructions added to correct a mistake.


2. Complete the following statement in not more than 10 words:

"In only 24 weeks, The Home Computer Course will make me..."

(The competition rules appeared in issue 1.)

This competition is open to readers in the UK and EIRE only.

.....ORDER YOURBINDERS NOW!



The normal price of each binder is just £3.95 inc. p&p. Below are two ways for you to buy them.

Week by week your copies are growing into the first complete guide to the world of home computers. They are specially written to help you make the most of your machine. Convert your copies into two course manuals with these attractive and specially designed binders.

Very special binder offer to save you £1!

A. If you buy Volumes 1 and 2 at the same time we will send them to you together with an Invoice for only £6.90 (inc. p&p) – saving you £1. Nothing could be easier – tick the top box on the order form opposite.

Or: B. If you prefer buying binder one at a time – please just send us your cheque/Postal Order for only £3.95 inc. p&p. We will send you your binder for Volume 1 only.

The Orbis Guarantee If you are not entirely satisfied with your binder you may return it to us within 14 days and any money you may have paid will be promptly refunded.

Important This offer is open only whilst stocks last. Please allow 28 days for delivery.

Overseas readers: This binder offer applies to readers in the UK, Eire and Australia only. Readers in Australia should complete the special loose insert in issue 1 and see additional binder information on the inside front cover. Readers in New Zealand and South Africa and some other countries can obtain binders now. For details please see inside the front cover.

Binders may be subject to import duty and/or local tax.